

APPLICATION FOR UNITED STATES PATENT

FRAUDULENT MESSAGE DETECTION

By Inventors:

Jonathan Oliver
3250 Ash Street
Palo Alto, CA 94306
A Citizen of Australia

Scott D. Eikenberry
3250 Ash Street
Palo Alto, CA 94306
A citizen of the United States

Gleb Budman
3250 Ash Street
Palo Alto, CA 94306
A citizen of the United States

Brian Kim
3250 Ash Street
Palo Alto, CA 94306
A citizen of Korea

Assignee: MailFrontier, Inc.

VAN PELT AND YI, LLP
10050 N. Foothill Blvd., Suite 200
Cupertino, CA 95014
Telephone (408) 973-2585

FRAUDULENT MESSAGE DETECTION

CROSS REFERENCE TO RELATED APPLICATIONS

This application claims priority to U.S. Provisional Patent Application No. 60/501,481 (Attorney Docket No. MAILP010+) entitled FRAUDULENT MESSAGE
5 DETECTION filed September 8, 2003, which is incorporated herein by reference for all purposes.

FIELD OF THE INVENTION

The present invention relates generally to electronic messages. More specifically, classifying messages is disclosed.

BACKGROUND OF THE INVENTION

10 There has been an increasing number of Internet scams designed to trick people into revealing personal information such as credit card numbers, passwords, social security numbers, etc. to unscrupulous individuals who intend to use the information for fraudulent purposes. These scams, also known as "phishing", usually involve sending
15 electronic messages such as e-mail to unsuspecting users. These messages are designed to appear as though they come from a reputable source, for example e-commerce companies, banks, credit card companies or other well-known people or institutions. The messages are often disguised as security alerts, account maintenance or renewal notices. The messages typically instruct the recipients to enter their personal information in the

message, or direct the recipients to a website operated by the fraud message sender where the information may be entered.

Since the messages typically appear to have a legitimate purpose, they often escape the detection of the conventional filtering software used to detect junk electronic messages (commonly referred to as "spam"). In some cases, even when the spam
5 filtering software detects that a message is junk and places it into a spam folder, because the message appears to be authentic, the recipient may reclassify the message and respond to it. It would be desirable to have a technique that would identify the fraudulent messages and prevent users from becoming victims of such scams.

BRIEF DESCRIPTION OF THE DRAWINGS

Various embodiments of the invention are disclosed in the following detailed description and the accompanying drawings.

Figure 1 is a system diagram illustrating a system used to detect fraud messages
5 according to one embodiment.

Figure 2 is a flowchart illustrating a message classification process according to one embodiment.

Figure 3 is a flowchart illustrating a message classification process according to another embodiment.

DETAILED DESCRIPTION

The invention can be implemented in numerous ways, including as a process, an apparatus, a system, a composition of matter, a computer readable medium such as a computer readable storage medium or a computer network wherein program instructions
5 are sent over optical or electronic communication links. In this specification, these implementations, or any other form that the invention may take, may be referred to as techniques. In general, the order of the steps of disclosed processes may be altered within the scope of the invention.

A detailed description of one or more embodiments of the invention is provided
10 below along with accompanying figures that illustrate the principles of the invention. The invention is described in connection with such embodiments, but the invention is not limited to any embodiment. The scope of the invention is limited only by the claims and the invention encompasses numerous alternatives, modifications and equivalents. Numerous specific details are set forth in the following description in order to provide a
15 thorough understanding of the invention. These details are provided for the purpose of example and invention may be practiced according to the claims without some or all of these specific details. For the purpose of clarity, technical material that is known in the technical fields related to the invention has not been described in detail so that the invention is not unnecessarily obscured.

20 A technique for classifying a message is disclosed. In some embodiments, the technique uses reference points in the message to classify the message. In some

embodiments, the message is determined to be a phish message if divergent reference points are detected. In some embodiments, fraud indicators are identified in the message and a statistical analysis is used to determine whether the message is fraudulent.

Figure 1 is a system diagram illustrating a system used to detect fraud messages according to one embodiment. Incoming messages such as message 100 are examined by fraud detection engine 102. The fraud detection engine may be a stand-alone software and/or hardware component, or integrated into a spam filtering system, mail system, or any other appropriate software code and/or hardware devices. It may be deployed on client devices such as personal computers as well as message gateways used to filter messages at an enterprise level. The fraud detection engine analyzes the message and determines whether it is a fraudulent message. Details of the detection process are discussed below. In some embodiments, once a fraudulent message is detected, a thumbprint or signature of the message is optionally computed. The thumbprint may be computed using techniques such as applying a hash function to the content of the message. The thumbprint may be stored in a database 104 to more efficiently identify this fraudulent message in the future. In some embodiments, the thumbprint database may be shared among users on a network.

To make the fraudulent message appear authentic, the body of this type of messages usually includes reference points to a legitimate source, such as URL links to logo images or web pages of well-known company, e-mail addresses with well-known domain name, as well as any other contact or identity information associated with well-known entities. As used herein, a reference point refers to information in the message

that can be used to contact the referenced entity or establish the identity of the referenced entity; a legitimate source refers to an entity that can be verified as being reputable. In order for the fraudulent message to fulfill its intended purpose, it should provide the recipient with ways to supply the desired information to the sender. Some of the

5 fraudulent messages include a reference point to a questionable source or entity whose reputation cannot be easily verified. The recipient is often instructed to provide his/her personal information to the questionable source by, for example, entering information into the message or by following a link to a website operated by the sender. This type of fraudulent messages, referred to as phish messages, usually include in the message body
10 at least one reference point to a source that is legitimate and at least one reference point to a questionable source. The characteristic of having a legitimate reference point and a questionable reference point is referred to as divergence, and such reference points are referred to as divergent reference points. Divergence can be used in some embodiments to identify phish messages.

15 Figure 2 is a flowchart illustrating a message classification process according to one embodiment. In this example, reference points are extracted from a message (200). In some phish messages, the reference point to the questionable source is disguised to trick the reader. For example, the URL
http://scgi.ebay.com@%32%31%31%2E%33%37%2E%31%38%35%2E%33%30:%37
20 %33%30%31 appears to be a link to "scgi.ebay.com" to the casual reader, but it actually links to a web page at a server with the IP address of 211.37.185.30, on port 7301. In

some embodiments, both the reference point to scgi.ebay.com and the reference point to the numbered server are extracted.

The reference points are then classified (202). In some embodiments, the reference points are looked up in a database that includes references of legitimate sources. There are various ways to establish and maintain the database. In some embodiments, users on a network can collaborate and vote on various reference points to determine whether they are legitimate, and the results of the collaborative effort are stored in the database; in some embodiments, collections of spam messages and good messages can be analyzed and the reference points in the messages can be extracted and stored to the database according to the message classification; in some embodiments, a registration process can be used to collect information from legitimate businesses and organizations; in some embodiments, existing databases of good reference points such as classified URLs, domains and phone number can be imported. In some embodiments, a reference point is classified as legitimate if the reference point or the source referenced can be found in the database; otherwise, the reference point is classified as unknown. In some embodiments, the reference point may also be looked up in a database of known offenders. A reference point found in this database may be classified as bad. Both unknown and bad reference points are classified as questionable reference points in some embodiments.

In the example shown in Figure 2, the message is then examined for divergent reference points (204). In some embodiments, different types of reference points are counted, and if there is at least one good reference and at least one questionable reference

in the message body, then the message is deemed to have divergent reference points. It is then determined whether the message has divergent reference points (206). A message that does not have divergent reference point is less likely to be a phish message and is processed as usual (208). In some embodiments, the message is delivered to the

5 recipient. In some embodiments, the message is further processed by spam classifiers. If, however, the message is determined to have divergent reference points, it is potentially a phish message and is handled appropriately (210). The handling of the message may vary for different implementations. In some embodiments, the message may be deleted or quarantined. In some embodiments, the recipient may be alerted about the phish
10 message. The alert can take many forms; for example, an indicator such as "[fraud]" may be added to the subject line, the recipient may be notified that a message has been quarantined or deleted, etc. In some embodiments, the recipient may be given detailed explanation of why the message is likely to be fraudulent. For example, both the legitimate source from which the message appears to be originated and the questionable
15 source designed to misguide the recipient may be clarified for the user.

In some embodiments, a statistical approach is used to classify phish messages as well as other types of fraud messages. Figure 3 is a flowchart illustrating a message classification process according to another embodiment. In this example, fraud indicators are used to classify fraudulent messages. Fraud indicators are patterns that are
20 characteristic to fraudulent messages. For example, having divergent reference points is one type of fraud indicator. More details of fraud indicators are discussed below. In some embodiments, the fraud indicator is identified by parsing the message and locating

certain patterns (300). A statistical analysis is then applied based on the fraud indicators (302). There are various statistical techniques applicable for the purposes of identifying fraud messages. In some embodiments, the number of fraud indicators is counted and the likelihood of the message being a fraudulent message is computed based on the number
5 of fraud indicators. Other statistical approaches for analyzing the indicators include neural networks, decision tree and classification tree techniques, logistic regression, as well as any other techniques or combination of techniques for analyzing features of the message according to a statistical model.

It is then determined whether the message is a fraudulent message (304). In some
10 embodiments, the message is determined to be fraudulent if the number of fraudulent indicators exceeds a threshold. In some embodiments, the statistical analysis computes a score that is used to determine whether the message is fraudulent. In some embodiments, the statistical analysis assigns weights to different indicators and computes a weighted score used to determine whether the message is fraudulent.

15 If the statistical analysis determines that the message is not a fraudulent message, it is processed as usual (306). In some embodiments, the message is delivered to the recipient. In some embodiments, the message is further tested using other spam filtering techniques. If the message is determined to be a fraudulent message, it is handled appropriately so that the recipient is prevented from responding to the fraudulent message
20 (308). The message may be deleted or quarantined in some embodiments; alerts and explanations may be sent to the recipient in some embodiments. In some embodiments, a

probability of being fraudulent is assigned to the message and the message is processed according to the probability.

The statistical message classification process shown in Figure 3 may be used in conjunction with the message classification process shown in Figure 2. In some

5 embodiments, once it is determined that the message has divergent reference points (206), a statistical technique is applied to further ascertain the nature of the message.

There are several fraud indicators that may be used by the statistical analysis, including raw Internet protocol (IP) addresses, non-standard encoding of URL's or addresses, links with user names embedded, misleading links, mismatched link names, message including
10 forms, forms with certain special requests, suspect content, etc.

Some fraudulent messages disguise the URL to the questionable source by using raw IP addresses such as http:// 211.37.185.30:7301. Therefore, raw IP addresses may be used as a type of fraud indicator in some embodiments. The identification of a raw IP address may be achieved by, for example, locating a string with a pattern of four numbers
15 separated by periods.

In some embodiments, non-standard encoding is used as a type of fraud indicator. Fraudulent messages sometimes use non-standard encoding of URL to obfuscate their addresses, such as the URL

http://scgi.ebay.com@%32%31%31%2E%33%37%2E%31%38%35%2E%33%30:%37
20 %33%30%31 shown in a previous example. There are a variety of ways for determining

whether non-standard encoding is used in a message, including finding % signs in URL links, using raw IP numbers in decimal, dword, octal or hexadecimal format, etc.

In some embodiments, URL links embedded with user names are used as a type of fraud indicator since fraudulent messages sometimes use this type of links. For example, `http://USER@netgo.port5.com` has a user name embedded in it. This type of fraud indicator may be identified by searching for an "@" symbol after "http://" or "https://".

In some embodiments, URL links embedded with misleading links are used as a type of fraud indicator since fraudulent messages sometimes include this type of URL. A misleading link disguises the questionable source by incorporating a reference to a well-known source. For example, `http://www.paypal.com@netgo.port5.com/index.php` includes a user name "www.paypal.com" that has the format of a domain name. Without closely inspecting the link, the recipient is likely to mistake www.paypal.com as the linked website rather than netgo.port5.com. This type of fraud indicator may be identified in some embodiments by searching for an "@" symbol followed by a domain name, where the text prior to the "@" symbol includes a substring such as ".com", ".net", "www", etc.

In some embodiments, mismatched link names are used as a type of fraud indicator. For example, a hyper text markup language (HTML) document with the URL ``
`http://www.BestBuy.com/fraud_department.html ` actually links to the domain

"www.digitalgamma.com"; however, the standard techniques for rendering HTML usually displays the link to the recipient as if it is a link to "www.BestBuy.com". This type of fraud indicator may be detected in some embodiments by comparing the actual link and the displayed link.

5 Forms in messages are used as a type of fraud indicator in some embodiments. Although forms are not commonly used in legitimate messages, fraudulent messages sometimes include forms in order to obtain information from the recipient. Thus, in some embodiments, the message is inspected for the <FORM action=...> tag that is embedded in the message to determine whether the fraud indicator is present.

10 Form requesting special information such as personal information is another type of fraud indicator that is used in some embodiments. The indicator is present if the message is determined to include a form, and the content of the form starting from the beginning "<FORM>" tag to the end "</FORM>" tag (or the end of the message if the end "</FORM>" tag is not present) includes requests for personal information. Whether
15 the form is requesting personal information may be determined by looking up key terms such as "Password", "Credit card", "Social security number", "Expiration date", "PIN", and any other terms pertaining to personal information.

 In some embodiments, suspect content is used as a type of fraud indicator. Suspect content may be detected by searching for certain key words or phrases such as
20 "verification of your account", "click on the link below to login", "fraud alert", "please confirm your account ownership by entering the information below", etc.

A technique for classifying a message has been disclosed. The technique is applicable to phish messages as well as other types of fraud messages. For example, the technique may also be applicable to messages used to commit brand fraud, where the recipient is shown a product of a certain brand, but directed to a website unrelated to the
5 brand-name product to purchase a different product.

Although the foregoing embodiments have been described in some detail for purposes of clarity of understanding, the invention is not limited to the details provided. There are many alternative ways of implementing the invention. The disclosed embodiments are illustrative and not restrictive.

10 WHAT IS CLAIMED IS: